PRE-APPEAL BRIEF REQUEST FOR REVIEW

Applicants incorporate and restate prior arguments made in response to rejections made in previous Office Actions. Claims 1-25 and 28-52 are pending in the application. Rejected claims 1 and 28 are independent claims. The Examiner has indicated that claims 5, 6, 8-10, 12-17, 19, 20, 22-25, 32, 33, 35-37, 39-44, 46, 47 and 49-52 would be allowable if rewritten into independent form. Applicants respectfully request review of the rejections of the remaining pending claims.

The Rejections of the Independent Claims Are Not Supported By a Prima Facie Case

Claims 1, 2, 3, 28, 29 and 30 are rejected under 35 U.S.C. 103(a) as being unpatentable over U.S. Patent No. 6,549,938 to Kilkki et al. ("Kilkki") in view of U.S. Patent No. 6,469,982 to Henrion et al. ("Henrion").

Claim 1 (and claim 28 in means plus function form) explicitly require the steps of, *inter alia*:

> receiving a stream of data from the switching fabric;
> extracting flow identity information from the stream;
> updating counters corresponding to the stream;
> subjecting the stream to a decision making algorithm in the bandwidth scheduler based on the extracted flow identity information and the updated counters for that particular stream resulting in that the stream is accepted or rejected before said stream enters any queue of said switch.

The Office Action fails to set forth a prima facie case of obviousness because neither Kilkki nor Henrion, alone or in combination, discloses or suggests at least: (1) deciding whether to accept or reject a stream based on updated counters for that particular stream before the stream enters any queue of the switch; and (2) deciding whether to accept or reject a stream based on extracted flow identity information for that particular stream.

The Advisory Action mailed April 17, 2006 also takes new positions that are clearly erroneous and which further fail to establish a prima facie case of obviousness.

In support of the rejections, the Office Action relies primarily on Figure 9 and the associated descriptions in the specification of Kilkki. Accordingly, Figure 9 of Kilkki is reproduced below for convenience.

FIG.9

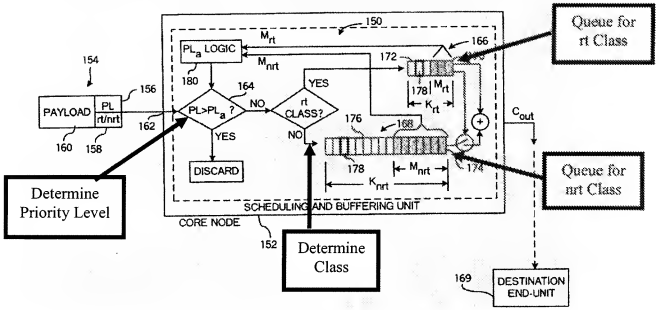

As can be clearly seen from Figure 9, Kilkki's "scheduling unit" compares a priority level (PL) of an ATM cell 154 to a threshold priority (PL_a) when deciding whether to discard the cell. It then places non-discarded cells in either a "rt" class queue 166 or a "nrt" class queue 168.

The Office Action also relies on Henrion, the figure for which is reproduced below.

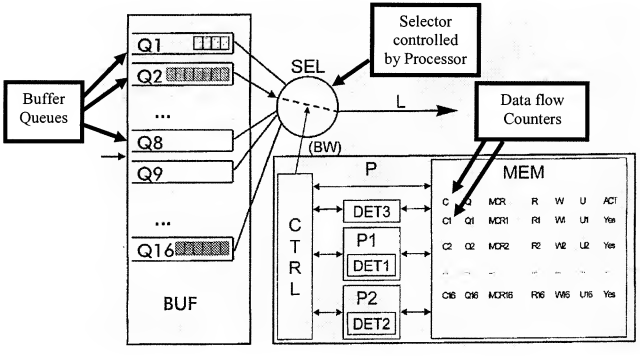

Henrion discloses an "intelligent buffer" that schedules flows of data out of a buffer BUF according to bandwidth reservation and sharing policies. As can be clearly seen above, data for respective flows are queued in queues Q1 – Q16, and then a processor P controls a selector SEL to allow the already-queued data to flow out of the buffers.

Neither Kilkki nor Henrion discloses or suggests deciding whether to accept or reject a stream based on updated counters for that particular stream before the stream enters any queue of the switch

The Advisory Action admits that Henrion's "flows are already in buffers Q1-Q16" before counters are updated and a decision is made whether to forward the flow. However, the Advisory Action further states that "the scheduler SEL must utilize the updated counters to determine which queue transmits a packet onto a common link L. The common link L reads on a queue of the switch. . . ." In other words, the Advisory Action posits that a decision is made before flows are sent to a forwarding link or "queue" L, but ignores that the counters are not updated and the decision is not made until after the flow data is already stored in the previous queues Q1-Q16.

Meanwhile, the claims explicitly require that the decision whether to accept or reject a stream based on updated counters for that particular stream must occur before the stream enters **any** queue of the switch. The Advisory Action's position is undercut by its own admission that the decision of scheduler SEL is made **after** the "flows are already in buffers Q1-Q16" (i.e. queues of the switch different from forwarding link L).

For at least this reason, the rejection of independent claims 1 and 28 is clearly in error and should be withdrawn.

Neither Kilkki nor Henrion discloses or suggests deciding whether to accept or reject a stream based on extracted flow identity information for that particular stream before the stream enters any queue of the switch

Applicants continue to disagree with the Office Action's position that Kilkki's priority level PL meets the required "extracted flow identity information."

The Advisory Action states that "the priority level reads on the flow identity since it distinguishes flows from each other." In other words the Advisory Action takes the position that

Kilkki distinguishes **flows** from each other. However, this ignores that a flow can be **stream** of many cells or packets all sharing a common identity (e.g. source/destination). Kilkki merely looks at PL levels of individual packets, it does not care or distinguish flows of many cells or packets.

Meanwhile, the claims clearly require extracting flow identity information from the stream and subjecting the stream to a decision making algorithm in the bandwidth scheduler based on the extracted flow identity information for that particular stream. In Kilkki's system, many packets can have the same priority level, but can belong to many different flows (i.e. connections between a particular source and destination). Contrary to the Advisory Action, Kilkki simply does not distinguish between different streams having different flow identity information as is explicitly required by the claims.

Kilkki itself clearly recognizes the difference between identity of flows or streams and priority information in headers of individual cells. Kilkki teaches that the priority level of the cell is simply a value that is placed in the cell by upstream nodes depending on upstream processing. Kilkki explicitly teaches that the scheduling unit in Figure 9 "make[s] decisions regarding the acceptance and rejection of packets based on the priority of a packet and the occupancy levels of the buffers in the particular core node 104. The core nodes, therefore do not need to obtain information regarding the properties of individual connections" (i.e. flows) (col. 12, lines 19-23, emphasis added). In other words, Kilkki explicitly recognizes that the properties of individual connections (i.e. flow identity information) are distinct and separate from mere packet priority information. Moreover, Kilkki teaches that upstream processors make priority decisions based on flow identity and teaches away from making decisions based on extracted flow information for a particular stream in "core nodes" such as the node that includes the scheduling unit in Figure 9.

Accordingly, Kilkki and Henrion does not disclose or suggest the explicit limitations of independent claims 1 and 28 for this additional reason, and the rejection thereof should be withdrawn.

Conclusion

For at least the foregoing reasons, independent claims 1 and 28 clearly define over Kilkki and Henrion. Accordingly, the 103 rejections of these claims, along with claims 2-4, 7, 11, 18, 21, 28-31, 34, 38, 45 and 48 that depend directly or indirectly therefrom, should be withdrawn. If any issues remain which the Examiner feels may be resolved through a telephone interview, s/he is kindly requested to contact the undersigned at the telephone number listed below.

Respectfully submitted,
PILLSBURY WINTHROP SHAW PITTMAN LLP

Date: May 31, 2006

40,580
Reg. No.

Mark J. Danielson
(650) 233-4777
Please reply to customer no. 27,498